Design and Evaluation of Inclusive Email Security Indicators for People with Visual Impairments

Yaman Yu*, Saidivya Ashok*, Smirity Kaushik, Yang Wang, Gang Wang
University of Illinois at Urbana-Champaign
{yamanyu2, sashok3, smirity2, yvw, gangw}@illinois.edu

***Abstract*—Due to the challenges to detect and filter phishing emails, it is inevitable that some phishing emails can still reach a user's inbox. As a result, email providers such as Gmail have implemented phishing warnings to help users to better recognize phishing attempts. Existing research has primarily focused on phishing warnings for sighted users and yet it is not well understood how people with visual impairments interact with phishing emails and warnings. In this paper, we worked with a group of users (N=41) with visual impairments to study the effectiveness of existing warnings and explore more inclusive designs (using Gmail warning designs as a baseline for comparison). We took a multipronged approach including an exploratory study (to understand the challenges faced by users), user-in-the-loop design and prototyping, and the main study (to assess the impact of design choices). Our results show that users with visual impairments often miss existing Gmail warnings because the current design (e.g., warning position, HTML tags used) does not match well with screen reader users' reading habits. The inconsistencies of the warnings (e.g., across the Standard and HTML view) also create obstacles to users. We show that an inclusive design (combining audio warning, shortcut key, and warning page overlay) can effectively increase the warning noticeability. Based on our results, we make a number of recommendations to email providers.**

1. Introduction

Globally, at least 2.2 billion people have some level of visual impairments [1]. Online services are often challenging for people with visual impairments to use. These users usually rely on assistive software such as screen readers to "read" the text line by line while moving their cursor over different parts of the web page. This process can be even more time-consuming and error-prone if the websites have accessibility or usability problems [2]–[4].

Recent research has explored the general security and privacy challenges faced by people with visual impairments [5]–[9]. Most studies are focused on *web browsing scenarios* such as performing login, and purchasing items online [10]–[12]. These studies have revealed a number of accessibility and usability issues such as poorly labeled login elements, inaccessible CAPTCHAs, audible password masking, and cumbersome password recovery mechanism for this user population.

In this paper, we focus on a critical scenario that has not been well understood by the existing literature, namely how people with visual impairments interact with and detect *phishing emails*. Since most email services today have implemented phishing warnings on their email interface, we particularly aim to understand the warning effectiveness on people with visual impairments. We seek to answer three main research questions:

- **RQ1:** Are users with visual impairments able to notice the security warning on phishing emails? If not, why?
- **RQ2:** Are users with visual impairments able to accurately detect phishing emails? If not, why?
- **RQ3:** How do users with visual impairments prefer email warnings to be presented?

To answer these questions, we worked with a group of n=41 users with visual impairments to conduct a series of studies (with IRB approval). Our work was primarily based on the Gmail platform, given that Gmail is the largest public email service [13] and is among the first few that implemented phishing warnings [14]. Our exploration had three key steps. In Step-1, we conducted a semi-structured interview study with 21 users to understand the challenges they faced when checking emails with screen readers. Our results showed that they were susceptible to phishing emails and the warnings were often missed. In Step-2, we worked with 5 (out of the above 21 users) to design a more *inclusive* email warning prototype to increase warning noticeability. In Step-3, we worked with another 20 users (who did not participate in Steps 1–2) to assess the inclusive warning design and compare that with the existing Gmail warnings across the standard view and the basic HTML view[1].

Findings. Our study has several important findings.

First, existing Gmail warnings were often missed because their designs (e.g., warning position, HTML tags used) did not match well with screen reader users' reading habits. For instance, most participants tended to use shortcut keys to quickly skip over various UI elements in the email head to reach the email body. As a result, warnings in the email head

*Co-first authors with equal contribution.

1. The basic HTML view has a simplified user interface (UI) and disables JavaScript-based features [15].

were often missed. Also, some Gmail warnings lacked the necessary HTML tags (e.g., `link`, heading), which made it difficult for screen readers to reach them.

Second, our inclusive warning design (prototype from Step 2) had higher noticeability than the current Gmail warnings and helped screen reader users detect more true phishing emails. The inclusive design combines audio warning (bell sound to draw user attention), warning overlay the email page (to make the warning easy to locate), and a shortcut key (for users to decide whether to expand the warning message).

Third, our study uncovered other issues in existing Gmail warnings such as the inconsistent warnings between the standard and HTML views. Such inconsistency existed even for *the same email* between different views. Another problem was there was no warning on the email list—we found that screen reader users often directly read the entire email on the email list (without opening an email page).

Contributions. This paper has three main contributions.

- We conducted interviews (n=21) to identify the challenges faced by users with visual impairments when interacting with phishing emails and warnings.
- We created inclusive designs to increase warning noticeability for people with visual impairments.
- We conducted another study (n=20) to assess the inclusive warning designs and offered a set of recommendations to email providers.

We have shared our findings with Google, and endeavor to connect with other email service providers (e.g., Yahoo, Outlook) to share the results. To facilitate future research, we also open-source our prototype designed in the paper[2].

2. Background and Related Work

Security Concerns of People w/ Visual Impairments. Researchers have studied the security and privacy concerns of people with visual impairments in both physical [5]–[7] and online contexts [8], [9], [12], [16]. Prior work has focused on web browsing scenarios such as web authentication [12], [17] and CAPTCHA solving [11]. Other studies have explored privacy concerns related to sharing images online [18] as well as using crowd-sourcing assistance services [19] and camera-based assistive technologies [9], [20]. Studies have shown that people with visual impairments are exposed to security/privacy risks when using online sources fraught with usability issues [10], [21].

Phishing Emails and Countermeasures. Phishing has been a persistent threat to Internet users. In a phishing attack, the attacker sets up *phishing websites* or sends *phishing emails* to lure the victim into giving away sensitive information. A large body of related work has been focused on phishing websites and their detection methods [22]–[32] and phishing URL indicators [33]–[36].

Our paper is more related to studies that are focused on *phishing emails* [37]–[45]. Unlike generic spam emails [46], phishing emails can be highly targeted and thus are more difficult to detect [47]. To deceive victims, attackers may spoof trusted entities as the sender address [14], [48].

As countermeasures, the most common approach is to detect and filter phishing emails before they can reach users [38]–[40]. However, this approach may not be able to remove all phishing emails. To this end, researchers also have studied ways to help users identify phishing emails in case they bypass the filtering and reach users' inbox.

Researchers have studied ways to improve users' ability to recognize phishing emails through training [49]. This can be done by improving users' understanding of phishing cues [50], [51] and their ability to parse phishing links [52].

Another approach is to implement security warnings on phishing emails. Recent work shows that it is beneficial to place the warning close to the suspicious link in the email [53]. Also active warnings (that require user action) are more effective than passive ones [53]. Other researchers have investigated the use of audio in email warnings [54]. Email warnings have been adopted by a number of real-world email services such as Gmail and Outlook [14].

Despite the extensive research efforts on phishing emails, most studies were not specifically addressing the challenges faced by users with visual impairments.

Phishing Threat to People with Visual Impairments. Only a few works have explored phishing threats to people with visual impairments. Focusing on *phishing websites*, researchers have found that it is more challenging for people with visual impairments to assess the credibility of a web page [4]. Due to the challenge to access the visual aesthetics and structural layout of a page, they rely on text and use a fast tab/scroll down the web page as an exploration tactic. Another study has evaluated browser extensions designed to protect users against phishing websites [55]. The results reveal a range of accessibility issues for people with visual impairments such as color-based security indications, missing instructions, and lack of shortcut keys.

A closely related work is the study from Blythe et al. [56] that focuses on *phishing emails*. The researchers interviewed eight people with visual impairments. They find that users with visual impairments are better at identifying phishing emails because they are more cautious and screen readers are helpful to capture grammar and spelling errors in phishing emails. Note that this study was conducted before email warnings were introduced to mainstream email services. Our paper advances existing literature by studying the efficacy of email warnings on people with visual impairments and exploring more inclusive warning designs.

3. Methodology Overview

We design and conducted a series of studies to explore the answers to our research questions. Fig. 1 provides an overview of our methodology, which contains three key steps. In step ❶, we perform an *exploratory study* to surface

2. https://github.com/yutouzai/inclusive_warning

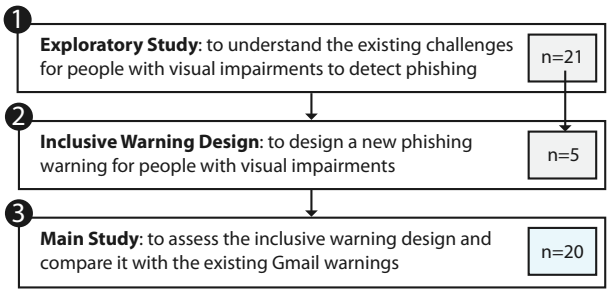

Figure 1: **Methodology Overview**—Note that participants of the main study (step ❸) do not overlap with those in steps ❶–❷.

the potential challenges for people with visual impairments to interact with phishing emails. One of the key observations from this study is that the security warnings on phishing emails are often not noticed by participants. This result further motivates our step ❷ where we explore new designs to make the phishing warnings more inclusive to people with visual impairments. In step ❸, we design the *main study* where we assess the inclusive designs in comparison with the existing designs in Gmail.

Our study focuses on Gmail as the experimental platform for two main reasons. First, Gmail is the largest public email service provider [13]. Second, Gmail is also among the first few services that implement phishing warnings [14]. As other email services often have similar warning designs, we use Gmail as a platform to explore general issues.

4. Exploratory Study

The first phase explored how people with visual impairments interact with phishing emails and email warnings (step ❶ of Fig. 1). We conducted an interview study with n=21 participants. Due to space limit, we only briefly describe the study procedure and discuss the key observations. Further details are presented in Appendix B.

During the study, participants played a role of an assistant to help their manager to process and review emails. Participants logged into a Gmail account set up by us, and the Gmail inbox contained three received emails (one legitimate and two phishing emails). They were instructed to read through the three emails and determine the action they would take for each one. Both phishing emails triggered the security warning in the Gmail standard view (the warning is shown as a graphical icon; see Fig. 11 in the Appendix). By default, participants were logged into the Gmail standard view. During the study, some participants preferred switching to the HTML view, and we allowed the switching.

We have two main observations from the study. First, the HTML view was preferred by a majority of the participants as 12 out of 21 participants asked to switch to the HTML view. However, after switching to the HTML view, we observed that there were no longer warnings on the two phishing emails. This indicates that the warnings between the HTML view and the standard view are not always consistent. We will further discuss this issue in Section 5.1. Second, participants had difficulty identifying certain phishing emails, and more importantly, they rarely noticed the email warnings. Among the nine participants that used the Gmail standard view, only two participants noticed at least one warning on their own (one participant noticed the warning only after they were asked about email legitimacy and went back to check the emails again). Based on our observation, most participants either failed to reach the warning or tabbed through the warning quickly without realizing that they had skipped it. We will further investigate and discuss the reasons for missing warnings in the main study in Section 7.

Observation 1: Users with visual impairments are susceptible to phishing emails; the graphical warning can be easily missed by them when using a screen reader.

5. Inclusive Design

Motivated by the above observations, we further explored to improve the *noticeability* of the warning and make it more inclusive for people with visual impairments (step ❷ of Fig. 1). Below, we first introduce existing warnings of Gmail and then describe the inclusive design prototype.

5.1. Existing Gmail Warnings

Gmail has different types of warnings and they look different between the standard view and the HTML view.

Standard View. The standard view has three types of warnings: icon warning, text warning in the yellow banner, and text warning in the red banner. Fig. 2 (a) shows an icon warning which displays a red question mark on the sender photo/icon. Hovering over the icon will pop up a message: "*Gmail couldn't verify [sender address] actually sent this message (and not a spammer).*" The banner warnings are shown in Fig. 2 (b) and (c), which appear at the beginning of the email body. The banner's color varies: the red color indicates a higher risk than the yellow color. The corresponding wording of the warning is also different (the warning text is transcribed in Figure 13 in the Appendix).

HTML View. Interestingly, even for the same email, the warning can look different when switching to the HTML view. First, as observed in the exploratory study, an email with an icon warning in the standard view may no longer have a warning on the HTML view. We performed additional tests with other phishing email content. We found that for certain emails, their icon warning would be converted to a text warning under the HTML view, as shown in Fig. 3 (a) and (d). The conversion happened to some emails but not all (which seemed to be dependent on the email content). Second, the banner warnings in the standard view are also converted to text warnings in the HTML view as shown in Fig. 3 (b) and (c). The warnings are shortened and highlighted in yellow color. All of the warnings in the HTML view are located at the *head* of the email, between the sender address and the "reply-to" address. Interestingly,

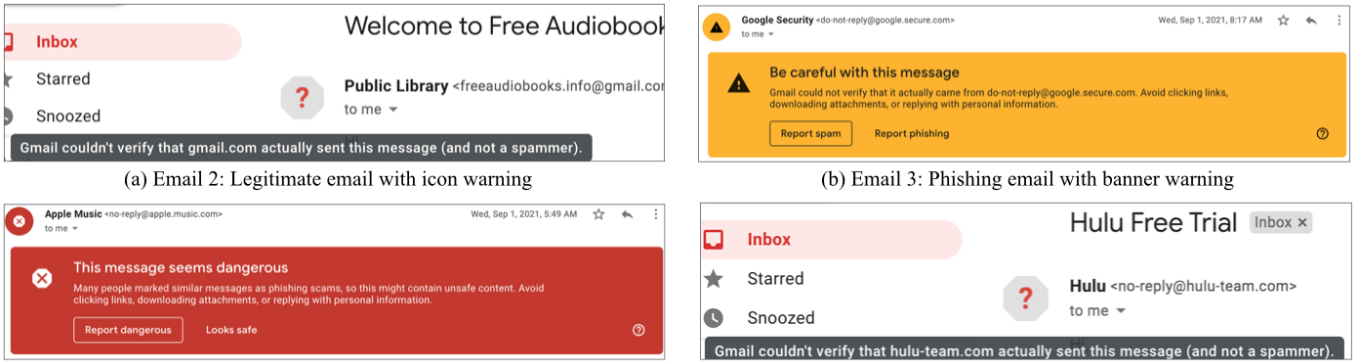

(a) Email 2: Legitimate email with icon warning

(b) Email 3: Phishing email with banner warning

(c) Email 4: Phishing email with banner warning

(d) Email 5: Phishing email with icon warning

Figure 2: **Gmail Warning in Standard View**—Emails to be used in the main study. Email 2 and 5 have the icon warning. Email 3 and 4 have the banner warning. Note that Email 2 represents a legitimate email with a misplaced warning.

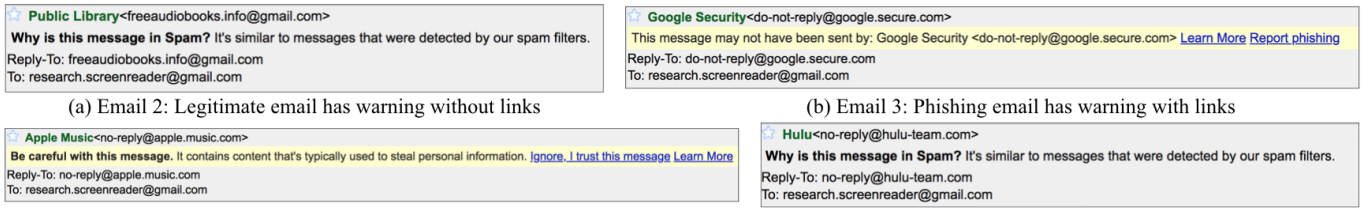

(a) Email 2: Legitimate email has warning without links

(b) Email 3: Phishing email has warning with links

(c) Email 4: Phishing email has warning with links

(d) Email 5: Phishing email has warning without links

Figure 3: **Gmail Warning in HTML View**—Emails to be used in the main study. In standard view (see Fig. 2), Email 2 and 5 have icon warning, and Email 3 and 4 have banner warning. Both types of warnings are displayed as text warnings in the HTML view.

the warning text is also changed between the standard and the HTML view. For example, the graphical icon warning is reworded as "*Why is this message in Spam...*" even though the emails are located in the inbox.

5.2. Inclusive Warning Prototyping

To make email warnings more inclusive (denoted as the "inclusive design") for screen reader users, we adopted a design probe [57] approach to explore the design space. We first developed the initial prototype of the inclusive warnings as a Chrome extension based on our observations from the exploratory study. We then contacted n=5 participants from the exploratory study to join a new session to provide feedback on the prototype and help us revise the designs.

5.2.1. Design Probes. Five participants were invited to join a new design session (30 minutes) with a $15 compensation each. Like the exploratory study, participants were instructed to read emails within a Gmail inbox and perform phishing detection, under their preferred interface (either the standard or HTML view). Unlike the exploratory study, we placed different types of warnings on these emails (four emails) and observed how participants interacted with them. Email 1 displayed the Gmail warning (red banner warning). The remaining three emails displayed different inclusive design options (detailed in Section 5.2.2). Specifically, Email 2 has an audio warning with a bell sound. Participants can use a shortcut key to expand the warning message, which overlays on top of the entire email page. Email 3 has an audio warning with a short speech (instead of a bell sound). Participants also use a shortcut key to expand the warning message. This time, the warning message only overlays on the email body not covering the email head. Email 4 does not have an audio warning. The warning message verlays directly on top of the entire email page. When participants were interacting with these warnings, we also orally explained the design and helped them to navigate. After completing all emails, participants were asked about their preference for different design options and their own designs. For example, "would you like to have audio security warning or not?" "Would you prefer to have the warning directly overlay the whole email page or have a shortcut key to control it?" "Do you have other suggestions to help screen reader users better notice email warnings? Please explain your design."

5.2.2. Key Design Choices. We now explain the key designs considered for the prototype and their differences with the existing Gmail design. Through the design probe sessions described above, we collected participant feedback and synthesized their inputs to design the final prototype.

Audio Warning. The first design choice is to use audio to draw users' attention. For sighted users, the current Gmail warning draws their attention with the red/yellow color, but the color can be less effective for people with visual impairments. Here, we consider using audio. While audio has been proposed for generic phishing warnings [54], there is a unique challenge for screen reader users: as screen readers continuously read content out loud, there is a need for the

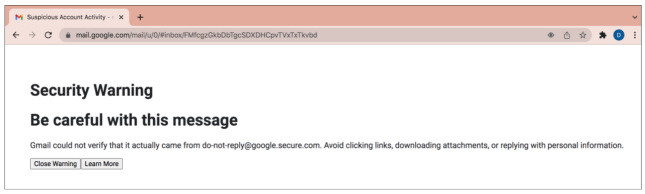

Figure 4: **Inclusive Design**—The warning first uses audio (bell sound) to draw users' attention, and prompts users to press a shortcut key (`control+1`) to expand the warning message. After the key pressing, the warning will overlay the email page. The content of the warning message is the same as that of Gmail.

audio warning to be clearly distinguished from the rest of the screen readers' speech. Based on the inputs from the pilot participants, we decided to use the *bell sound*, followed by a short introduction speech "*warning! for details, press control plus one.*" The bell sound is used to distinguish the warning from the rest of the screen reader's speech. It plays when the user opens an email that carries a warning.

Shortcut Key. After the audio warning, we prompt users to press a shortcut key `control+1` to expand the warning message[3]. There are two main considerations for including a shortcut key. First, screen reader users are already heavily using shortcut keys to navigate the web pages (based on our observations from the exploratory study). They are familiar with this mechanism. Second, the shortcut key gives users the choice of whether to expand the warning. Our pilots also confirmed users' preference to have such a choice.

Warning Overlay. The last question is how to present the warning text to users. Our design choice is to overlay the whole page after the user presses the shortcut key `control+1`, as shown in Fig. 4. This is motivated by a few reasons. First, our exploratory study shows that warning text can be easily missed by screen readers when it is mixed with other UI elements. By overlaying the warning on top of the email page, screen readers can easily locate the warning text. Second, prior research shows warnings that actively "interrupt" users' browsing process is more effective than passive warnings [33]. Users can resume email reading by clicking on the "close warning" button.

Note that these design choices were synthesized based on the feedback from five pilot participants. Therefore, we do not expect the prototype to be the optimal design for all users. Instead, we use this prototype in our main user study to explore the preferences of a larger number of users and inform further improvements (see Section 6).

Our goal is to improve the accessibility and noticeability of the warning. The specific text/wording of the warning message is not our focus, and we use the same wording of the Gmail warning for our experiment. In practice, email services can customize their own warning text and use the inclusive design to display it to users. Phishing detection algorithm is also out of our scope. It is up to the email providers to improve their detection algorithm to decide when to place the warning on suspicious emails.

3. We chose `control+1` to differ from existing shortcut keys in screen readers (e.g., JAWS, NVDA). Email providers can make their own choice and add the reminder of the shortcut key to each warning message.

6. Main Study Protocol

Using the above prototype, we conducted a user study to assess the inclusive design and compared it with existing Gmail warnings (step ❸ of Fig. 1).

6.1. Study Design

We recruited a new group of n=20 participants for the main study *who did not participate in the exploratory study*. During the interview session, participants were instructed to review emails under the same scenario used in the exploratory study (see Appendix A). We did not prompt the participants that the study was about phishing warnings. Instead, they played the role of an assistant to help their manager to process and review emails. Role-playing is a commonly used method to study phishing susceptibility [52], [58]–[60]. In the main study, participants would read five emails under our inclusive warning design and then read the same five emails under the Gmail design (within-subject). When reviewing the emails, the participants were asked to orally describe how they would process each email. In the end, we asked participants questions about the reasons behind their actions and their preferences on the alternative design choices for the warning.

Email Selection. We selected 5 emails (3 phishing, 2 legitimate) to cover different warning types. The three phishing emails were created based on examples found in the `MillerSmiles.co.uk` scam archive [61] and they were spoofing three popular email senders including Google, Apple, and Hulu. The two legitimate emails were selected from the authors' email inboxes and the messages were sent from a residential community and a public library. For further details of the emails, see Figure 13 in the Appendix.

Under the Gmail warning design, the three phishing emails (Email 3–5) had different types of warnings, as shown in Fig. 2 (b–d). For the two legitimate emails (Email 1–2), we added an icon warning to Email 2 (Fig. 2(a)) to emulate a false alert. Gmail's icon warning may show up on legitimate emails, for instance, if the benign email sender has misconfigured their SPF or DMARC protocols. The legitimate Email 1 did not have a warning (not shown in the figure). In the main study, these emails were carefully selected and configured, to make sure the warnings would consistently appear on the HTML view too (see Fig. 3).

The inclusive design condition had the same setup, namely, Email 2–5 had warnings and Email 1 did not.

Participant Recruitment. We recruited participants from social media, mailing lists, and snowball sampling (i.e., participants introduced friends to join our study). Interested users first took a screening survey where we collected basic information including age group, occupation, self-reported visual abilities, and their regularly used email services, browsers, and screen readers. Then we identified 20 eligible

participants to participate in our interview session. Eligible participants are those who (1) have visual impairments and (2) regularly use screen readers, Gmail, and Chrome browser in their daily life. The goal was to ensure participants were familiar with the environment we provided. The interview took 1.5-2 hours (including the setup time). Each participant was compensated $30. We also compensated participants $10 if they refer a new participant to join the study.

Interview Procedure. For eligible participants, we first provided an online consent form by email, informing our study procedure and data protection policy. Participants were informed that this study was designed to improve the accessibility of screen readers at the beginning to avoid priming. We later debriefed our participants about the true purpose of the study after they finished the email reading tasks.

We set up a virtual machine (VM) with the latest Chrome browser installed. We also pre-installed the screen reader based on participants' preferences (according to the screening survey). The 5 emails were placed in a Gmail account set up by us on the virtual machine. During the interview, we shared the screen of the VM and participants used the *remote control* function of Zoom to perform tasks on the VM. The remote control function allowed participants to control the VM as their own devices using mouse and/or keyboard. The interview was recorded with the participants' consent.

We used a virtual machine instead of using participants' own devices for two considerations. The first reason was to shorten the setup time of the interview. Using the participants' own devices would require them to download the inclusive design prototype (a browser extension) and log in to our Gmail account using a shared password. This process turned out to be extremely challenging for people with visual impairments. Based on our pilot test, the setup time could take 30–60 minutes. The second reason was to protect participants' privacy (they did not need to share the screen of their own devices).

During the interview, participants first configured the screen reader to ensure the reading environment was the same as their own device. Then participants read the same five emails under the Gmail design and also the inclusive design. We did not mention which design was the new design. Instead, we referred to them as the first design and the second design to avoid biasing participants' preferences. Using the HTML view or the standard view was a choice of the participants based on their own reading habits.

We only provided the minimum instruction at the beginning of the study. For the inclusive design, we informed participants that "*please read the five emails following your daily reading style. Please note that some of these emails might come with a warning. When you encounter a warning message, please press control+1 to listen to the warning.*" For the Gmail design, the instruction was: "*please read the five emails following your daily reading style. Please note that some of these emails might come with a warning. When you encounter a warning message, please listen to the warning.*". Note that we did not include the instruction of "*press control+1 to listen to the warning*" for the first 7 participants. Some participants mentioned that they were not familiar with the new shortcut key to read warning messages and asked us to provide a clear instruction. For later participants, we included this instruction by default.

After participants read each email, we asked them to verbally describe the actions they would take (they did not need to actually perform the action on the email). When participants finished reading five emails in one design, we asked them how many emails contained the warning and how many emails they considered as phishing. Then participants switched to the other design and reviewed the five emails again. Furthermore, after finishing all the email reading tasks, we showed participants additional alternative designs with different combinations of design choices (e.g., whether to use audio, whether to use shortcut keys, different ways to present the warning message), to understand participants' preferences. At the end of the study, we asked questions about their preferences of designs and their prior experience with phishing emails and security warnings.

6.2. Data Collection and Analysis

Our analysis was focused on the following types of data. First, we had user behavior data, namely, the recorded user actions on the email pages (e.g., moving the cursor over different UI elements, pausing the cursor to listen to the screen reader). Second, we collected participants' responses to our questions on how they would take action on the emails and their reasons. Third, we counted the number of warnings and phishing emails they noticed, and their preferences for different designs.

It was challenging to recruit a large number of participants from the user population with visual impairments. Therefore, we used a *within-subjects* design to collect data from n=20 participants where each participant read emails with both the Gmail design and the inclusive design. More specifically, nine participants read emails with the Gmail design first and the other 11 participants read emails with the inclusive design first. To mitigate the impact of repeated measures (i.e., each participant rated the same email message twice, each time under a different warning design), we only used each participant's *first-round data* when *quantitatively* comparing the two designs. We used both rounds of data in our *qualitative* analysis to understand the reasons why users noticed or missed warnings on the emails and how it affected their phishing detection.

We transcribed interview recordings and analyzed the qualitative data using thematic analysis [62]. Two co-authors coded the data and converged on a codebook with more than 40 codes, which were then categorized into three high-level themes: "warning noticeability," "phishing detection," and "design preferences." Details about this qualitative analysis are presented in Appendix D. When representing quotes from participants, we will mark whether they are from the first round or the second round.

7. Results

In this section, we present our results from the main study. We first describe participant demographics, followed by the analysis results to answer the three research questions regarding warning noticeability (Q1), phishing detection performance (Q2), and warning design preferences (Q3).

7.1. Participant Background

Our participant demographics are summarized in Table 3 in the Appendix. Participants were from different age groups: five reported in the 18–24 age group, nine in 25–34, four in 35–44, one in 45–54, and one in 55–64. Among the 20 participants, five participants (P3, P4, P7, P8 and P9) had low vision, and the other 15 were blind. Seventeen participants reported they are male and three reported being female. Our participants had diverse professional backgrounds, including one administrative staff, three business managers, three teachers, three students, one medical worker, one in economics, one self-employed, six computer engineers, and one engineer in other fields.

7.2. Warning Noticeability (Q1)

We first examine the noticeability of warnings for the Gmail design and the inclusive design (Q1). Below, we present the descriptive statistics of participants' actions, followed by regression analyses. Then we explore why participants missed warnings.

7.2.1. Overall Success Rate of Noticing Warning. As discussed in Section 6.2, we focused on the first-round results when statistically comparing the two designs. We observed that the inclusive design helped participants notice warnings in the emails (Fig. 5). The success rate of noticing warnings was based on the observation of the participants' actions on emails and their responses to our questions. During the email task, as the participants moved their cursor on the email page, the screen reader would read the corresponding text. This allowed us to observe whether the warning message was "read" to the participants by the screen reader. Furthermore, we also asked the participants what actions they would take for each email and whether they noticed warnings after they read all 5 emails. Based on their answers, we can determine the number of warnings noticed (out of 4 warnings). As shown in Fig. 5, no participants noticed all 4 email warnings under the Gmail design. Under the inclusive design, 9 out of 11 participants noticed all 4 warnings.

7.2.2. Regression Analysis. Even though the dataset is small, we run a linear mixed-effect model to determine the impact of inclusive design on warning noticeability. Using only the first round data, we treat each email review as a data point. The dependent variable is a binary variable, indicating whether the participant noticed the warning or not. The type of warning (Gmail or inclusive) and the Gmail view (Standard or HTML) are the independent variables.

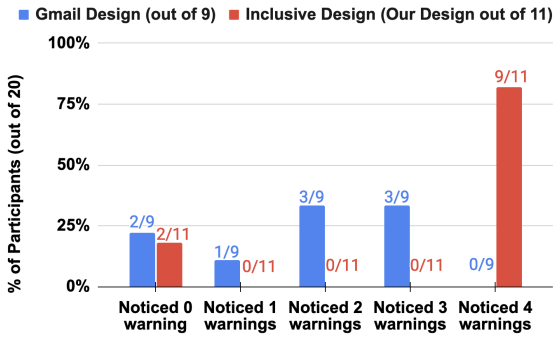

Figure 5: **Email Noticeability**—Success rate of noticing email warnings under the Gmail design and the inclusive design for participants reviewing them in first round of interviews. Our inclusive design helps people with visual impairments notice warnings.

Since each participant read the same five emails in both conditions, participant id and email id are random effects in the model. This analysis suggests that the inclusive design significantly increased warning noticeability (estimated coefficient = 0.3856, $p<0.015^{*}$, $r2 = 0.837$).

Observation 2: The inclusive design has a higher noticeability than the current Gmail designs.

7.2.3. Inclusive Design—Reasons for Missing Warnings. Participants rarely missed the warning under the inclusive design (2 out of 20). The two participants (P6, P7) both noticed the warning (inclusive design, first round) but did not press the shortcut key `ctrl+1` to expand the warning message. The reason was that they were not familiar with the new design. P6 stated *"control plus one, yes (I heard it). no (I didn't press), it is probably because it is your screen and I am not sure if I should be checking it, you know, I didn't want anything to go wrong for your screen."* P7 also claimed *"Yes, I heard the warning since like the second email. The security notification on this email is you press control plus one to get more information. (I am) not familiar with this kind of warning system."* As mentioned in Section 6, we added the instruction to avoid such confusion for later interviews, and all participants noticed and checked the warning messages.

7.2.4. Gmail Design—Reasons for Missing Warnings. The reasons for missing Gmail warnings are more complex. We can characterize them in three aspects: (1) the reading habit of screen reader users; (2) the position/location of the warning message; and (3) the HTML tags used to implement the warning. These factors (especially factors 2 and 3) have different impacts under the HTML view vs. the Standard view. Next, we will first describe the common email reading patterns with screen readers and then dive into the reasons for the HTML view and the standard view, respectively.

Patterns of Email Reading Habits. We identified four patterns by observing how participants interact with the different elements of an email and discussing the observed

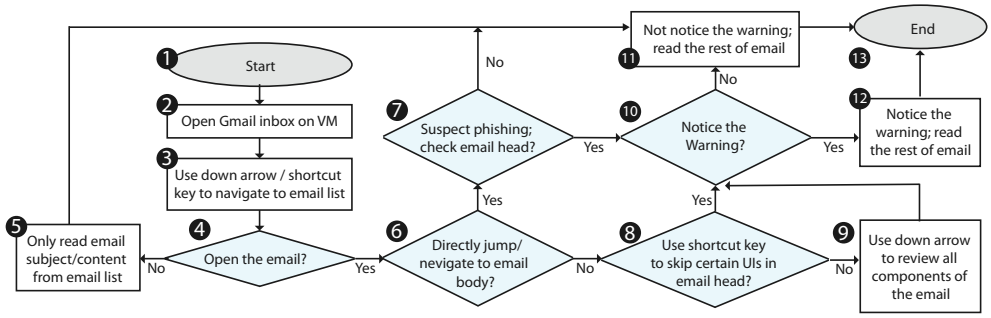

Figure 6: **Email Reading Patterns**—Flowchart of the email reading process of screen reader users.

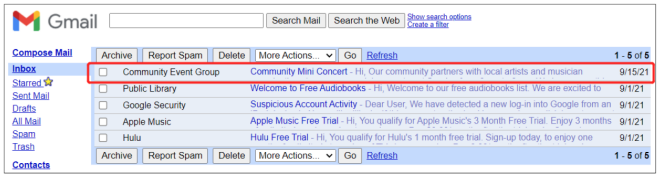

Figure 7: **Email List of Inbox**—A screen reader can read the email subject, sender name, and the email body directly from the inbox email list (for both Gmail HTML and Standard views).

behaviors with them. Note that one participant may have multiple patterns across different emails.

The first pattern is visualized by a sequence of steps (1-2-3-4-5) in Fig. 6. It represents a participant reading emails directly from the email list of the inbox. As shown in Fig. 7, the email list shows the sender name and email subject for each email. More importantly, a screen reader can also read the *entire email body* without opening the email on a new page (by pressing the down arrow key). During our study, only one participant P4 had this pattern explicitly. P4 clicked on the email but also waited for the screen reader to finish reading the entire email body from the email list. As a result, P4 did not read the opened email page and thus missed the Gmail warning. P4 noticed our inclusive warning because the audio warning was played right after the email was clicked on. However, if the participants did not click on the email (and only read from the list), they would have missed our warning too. Other participants, such as P15, also mentioned this habit (e.g., checking emails from the email list) in the exit interview (see Section 7.4). This suggests the security warning should also show up on the email list.

The second pattern corresponds to the sequence (1-2-3-4-6-7-*) in Fig. 6. The asterisk (*) here is a wildcard. Unlike the first pattern, participants opened the email page to read the email. Under this pattern, they directly jumped to the email body without reading the email head. This could be done using certain shortcut keys or by quickly pressing the down arrow. For example, P19 used shortcut keys to skip the email head when reading Email 2. This made P19 miss the warning in the email head. 8 participants had this pattern.

The third pattern corresponds to the sequence (1-2-3-4-6-8-9-*) in Fig. 6. Instead of directly jumping to the email body (the second pattern), participants started from the head of the email and used the down arrow key to read through the rest of the email page. For example, P8 explained: *"So I would press 'h', till the time I land upon that heading from where the email starts, then I press down arrow to read."* Under this pattern, if participants pressed the down arrow key to read through the email quickly, they still had the chance to miss the warning. 8 participants had this pattern.

The fourth pattern corresponds to the sequence (1-2-3-4-6-8-10-*) in Fig. 6. Participants also started from the head of the email. Instead of reading through the email page with the down arrow key (the third pattern), the participants used shortcut keys (`n` and/or `tab`) to skip over some UI elements in the email head to reach the email body. Different from the second pattern, the participants still let the screen reader read some UI elements (instead of skipping all). Based on the follow-up interviews, we found that participants were trying to skip UI components such as "reply," "email label," "print," and "more." These UI elements are often recognized as *links* by the screen reader, and thus the shortcut keys used are `n` (skip past link) and `tab` (next link). Since some Gmail warnings (e.g., icons) contain links and some are not (e.g., text), users may miss different types of warnings depending on the shortcut keys used. 8 participants had this pattern.

Observation 3: The current Gmail warning designs do not match well with screen reader users' reading habits (e.g., reading emails from the email list, skipping email head to jump to email body), and thus can be missed.

Missing Warnings in Gmail HTML View. For the HTML view, the reason for missing the Gmail warning is that the warning position and the HTML tags used do not match with users' reading habits. Also, the inconsistent warning designs further caused confusion for users.

First, the *position* of the warning matters—all Gmail warnings in the HTML view are located in the email head (see Fig. 3). Many participants would skip the head and jump to the email body using shortcut keys. Even if the participants read the email from the head and used the down arrow key to visit each element, they could still miss the warning when pressing the key quickly. For example, P14 pressed the down arrow key quickly and skipped the warning during the study. They stated *"that is possible because you read it fast. And you go further. It may so happen that you skipped off the warning."*

Second, the HTML `link` tags used in the warning also

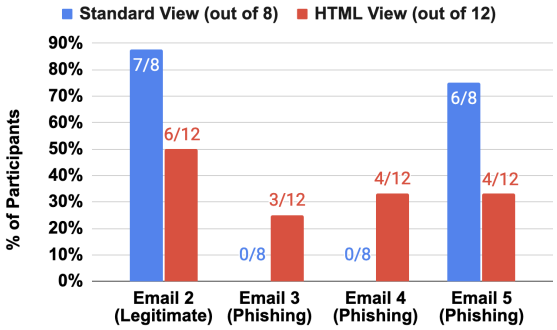

Figure 8: **Missing the Warnings (Gmail Design)**—The % of participants that missed the warning on each email (including the results of both rounds).

have an impact. The reason is related to the fourth reading pattern as participants use shortcut keys to skip certain UI elements such as the "reply" button (recognized as a `link`). For example, the Gmail warning on Email 2 is easily missed under the HTML view because the warning *does not* contain a `link` tag. When participants press `n` (skip past link) or `tab` (next link) to jump over parts of the elements before the email body, screen readers will directly skip the warning. Fig. 8, for the 12 participants who selected the HTML view, 50% (6/12) missed the warning in Email 2.

Third, the inconsistent warnings also caused confusion to participants. Recall that the warnings of Email 2 and 5 do not have a link while the warnings of Email 3 and 4 have links (Fig. 3). When participants used *the same* shortcut keys to read five emails, they encountered different results. For example, P19 always used `x` (reaching the checkbox), `shift+h` (reaching the previous heading), and `n` (skip past link) to reach the email body. P19 missed the warning in Email 2 in the head with these shortcut keys (Gmail design, first round). However, when reading Email 3 and 4 (Gmail design, first round), the same shortcut keys took P19 to the reply-to address (instead of the email body) due to the link tag. This surprised P19 and prompted him to check the warning: *"Something that I didn't notice in the Google Security email is that when I pressed n to jump to the body, it actually jumped to the warning. So here (Email 4) we have the same ... The warning is different (from Email 2)."*

Observation 4: Under the HTML view, Gmail warnings are missed primarily because they are all located in the email head. Some warnings are skipped by screen readers because they do not contain a `link` tag.

Missing Warnings in Gmail Standard View. For Gmail standard view, the *position* of the warning is still an important factor. However, the *link* tag is less so since all the warnings in the standard view contain link tags. The inconsistency issue still exists in the Gmail warning designs.

First, the position of the warning still matters. Recall that there are two types of warnings on Gmail's standard view (see Fig. 2). The icon warning (Email 2 and 5) is located in the email head and the banner warning (Email 3 and 4) is located in the email body. As shown in Fig. 8, there were eight participants who selected the Gmail Standard view. 88% (7/8) of them missed the warning on Email 2 and 75% (6/8) missed the warning on Email 5 which are both icon warnings. In comparison, no screen reader user missed the banner warning on Email 3 and Email 4. The result indicates that the banner warning (in the email body) matches better with screen reader users' reading habits than the icon warning (in the email head).

Second, while both icon warning and banner warning contain the `link` tags, they worked differently with screen reader users. Icon warning is located in the email head, and thus it can be easily missed as participants use `tab` (next link) or `n` (skip past link) to quickly navigate through the link components in the email head. Also, the icon warning is not a heading element, and can be missed by participants who use `h` (next heading) and `n` (skip past link). In comparison, the banner warning is already hard to miss as it sits at the beginning of the email body. The banner warning also contains link and heading tags, which can be easily reached by different shortcut keys.

Third, the inconsistency of warning designs is still a potential issue. In addition to the warning position and the HTML tags used, participants also mentioned other differences between the icon and banner warnings that affected the noticeability. For example, P10 pointed out that the banner warning's message is much longer, which is more noticeable: *"because there was quite a lot of warning text there for the Apple one [with a banner warning]."*

Observation 5: Under the standard view, the Gmail icon warning is often missed since it is located in the email head. Banner warning (located in the email body) is often correctly noticed by screen reader users.

7.3. Phishing Detection (Q2)

Next, we examine participants' phishing detection performance. For each email, participants described orally their assessment of the email and how they would take action. We categorize participants' responses into "Legitimate," "Suspicious," and "Phishing". We categorize a response as "legitimate" if the participant explicitly mentioned that they thought the email was genuine/legitimate or if they were not hesitant to reply or click on the link in the email. We categorize a response as "phishing" if they explicitly mention so. Sometimes, participants expressed their suspicion of the email but stated that they would need to perform extra steps to confirm their conclusion. For example, P15 stated that "*the URL of Email 4 should be iTunes. It should not be apple.music.com. So I will check that page to see what information I get from that page.*" In these cases, we categorize the response as "suspicious." For our analysis, we optimistically consider both "suspicious" and "phishing" as successful phishing identification.

7.3.1. Regression Analysis. We did not observe a statistically significant impact of warning design on phishing

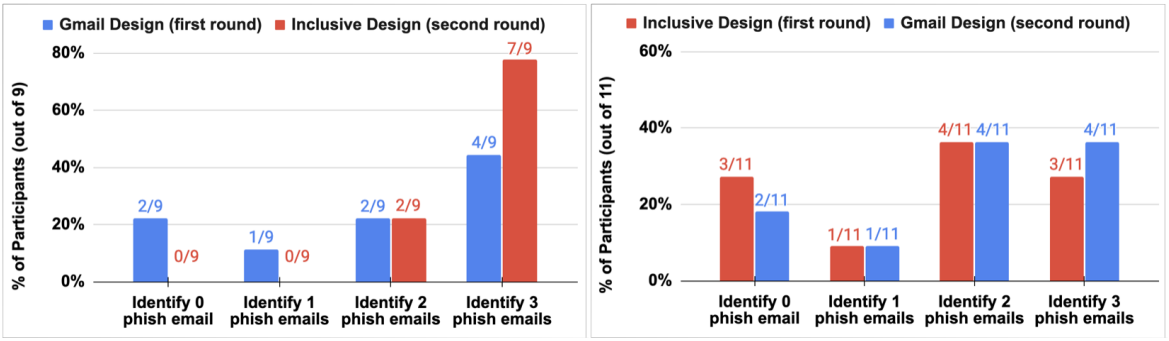

(a) Phishing detection with Gmail design first

(b) Phishing detection with Inclusive design first

Figure 9: **True Positives**—Number of participants detected true phishing emails under the Gmail design and the inclusive design for participants who review each design in the first round. A higher percentage of participants were changing their decision for the legitimacy of emails in the second round email reading when they experience Gmail design first then inclusive design.

detection outcomes (possibly due to the small data size). We use the same linear mixed-effect model as in Section 7.2.2 on the first-round data only. The only change is that the dependent variable is a boolean variable indicating whether a participant correctly labeled an email. We did not find significant evidence to suggest that the inclusive design impacted participants' phishing email detection (p>.05, $r2$ = 0.275). In the following, we instead focus on *qualitative analysis* and descriptive statistics to understand participants' decision-making and reasons behind detection errors.

Email	Condition	# Users detected phishing / # Users	
		Gmail Design	Inclusive Design
Email 2	Noticed Warning	0/1	3/9
	Not Notice Warning	0/8	0/2
	Total	0/9	3/11
Email 1	Total	0/9	0/11

TABLE 1: **False Positives**—Email 1 and Email 2 are legitimate emails. Email 2 has a misplaced warning. We report the number of participants that identified the email as phishing (over the total number of participants in each condition) in the first round.

7.3.2. Phishing Detection Results. We break down the results into true/false positives and true/false negatives. Due to space limits, we focus our discussion here on true positives and false positives. The discussion of true negatives and false negatives is presented in Appendix C.

True Positives. True positives refer to the true phishing emails being correctly identified.

Fig. 9 shows the number of true phishing emails identified by participants who reviewed emails in the Gmail design first (Fig. 9(a)) and the inclusive design first (Fig. 9(b)), respectively. Recall that we have three true phishing emails (Email 3, 4, and 5). As shown in Fig. 9(a), participants who reviewed the Gmail design first tended to change their decision in the second round when viewing emails with the inclusive design (i.e., identified more phishing emails in the second round). For example, 64% of participants identified all three phishing emails during the second round (inclusive design) compared that of 44% in the first round (Gmail design). We did not observe a similar trend for participants who reviewed the inclusive design first.

The results indicate that warnings may have changed participants' perceptions of an email. For example, P11 did not notice the warning of Email 5 under the Gmail design in the first round and stated that *"This one (Email 5) does not have that warning. So it is a promotional email"*. This participant changed his opinion when he noticed the warning of Email 5 under the inclusive design in the second round: *"a suspicious email warning on that email (Email 5), so I'll ignore that (email)."* Similarly, by noticing more email warnings with the inclusive design (second round), P4, P10, P12 and P13 also identified more true phishing emails.

However, improved noticeability of warnings may not always improve phishing detection rate. This is because some participants made their decisions based on other factors (e.g., email content). For example, P16 noticed all warnings of emails with inclusive design in the first round but perceived all phishing emails as legitimate, because the participant relied on the content to judge if an email should be trusted. P14 also considered the two phishing emails as legitimate with the inclusive design in the first round, because she paid attention to the (spoofed) email sender address and thought they "*looked familiar*." Some participants misinterpreted the warning messages. For instance, P8 thought the warning of Email 5 (standard view) was stating that the email sender is not Gmail but another email service. The participant perceived that warning as normal (not harmful). P13 had a similar misunderstanding: *"because the security warning (of Email 5) does not seem too harmful like before. I would like to read the content of the email. I will take the free trial from the URL in the email."*

False Positives. False positives refer to legitimate emails that were incorrectly identified as phishing. Recall that we have two legitimate emails (Email 1 and Email 2) and we have intentionally added a warning to Email 2 to emulate misplaced warnings. Overall, the result shows that participants are more likely to falsely identify a legitimate email as phishing if the email carries a misplaced warning. As shown in Table 1 (first-round result), under the inclusive design,

Different Design Combinations	# of Participants	Percentage
Audio warning + Security warning on the same page with email body	9	45%
Audio warning + Key press + Security warning on the same page with email body	4	20%
Audio warning + Security warning replacing the email body	3	15%
Audio warning + Security warning overlay the whole page	2	10%
Security warning on the same page with email body	2	10%

TABLE 2: **Design Preferences (Combination)**—Preferred design choice combinations among participants.

Email 2 has more false positives (3 out of 11 participants) compared with Email 1 (none of 11 participants). This is likely due to the misplaced warning on Email 2. In Table 1, we also break down the Email 2 results based on whether the participants noticed the warnings (first two rows). The result shows that all of the false positives were made under the condition that the participants *noticed* the warning. The inclusive design had a few more false positives (3 out of 9 participants) than the Gmail design (none of 1 participant) because more participants noticed the warnings under the inclusive design.

We argue that the warning is supposed to be easily noticed by users. It is the email providers' responsibility to avoid misplacing warnings on legitimate emails. In Gmail's case, the icon warning is often misplaced due to the misconfigured SPF/DMARC protocols of the *email senders* [14]. As such, the false positives can act as a forcing function for legitimate email senders to resolve such misconfigurations.

Observation 6: While noticing warnings helped some participants in identifying phishing emails, participants also examined other factors (e.g., email content) when making assessment decisions. Meanwhile, misplacing warnings on legitimate emails would lead to a higher rate of false detection regardless of warning design choices.

7.4. Preference of Warning Design (Q3)

So far, we have shown that our inclusive design prototype can improve warning noticeability and in turn phishing detection efficacy. However, we do not expect this is the optimal design for everyone. After the email reading task, we asked participants about their preferences for each design feature. These features include audio warning, keypress (i.e., pressing a shortcut key to show the warning) and three different ways to present the warning. The choices for presenting the warning include (1) showing email warning on the same page with the email body, (2) email warning replacing the email body, and (3) email warning overlaying the whole page. Note that (2) and (3) require participants to click on the "close" button to resume reading the original email content. Participants answered whether they preferred to include the feature or not.

The most preferred features are using audio warning (18/20=90%) and presenting the warning on the same page with the email body (15/20=75%). Table 2 further illustrates the popular design combinations. The results suggest that participants' preferences are diverse. 9 out of 20 (45%) participants prefer the audio warning and warning on the same page with the email body. An additional 20% (4/20) prefer to also include the shortcut keys in this design. Interestingly, "warning overlay over the page" was not the commonly preferred choice, despite it being well received during the prototyping phase (Section 5.2). Recall that the overlay design intentionally disrupts the email reading process to draw users' attention to the warning. Sometimes, the interruption can be perceived as excessive. For example, P2 explained *"when I clicked a thing, it overlays the email ... it is taking too much control from me."* We believe both "warning overlay" and "warning in the email body" are reasonable designs given their improved noticeability over "warning in the email head" (see Section 7.2). Further trade-offs can be made between disruption to users and the risk level of the email (e.g., using overlay for only high-risk emails).

To further probe the participants, we asked them about their *desired features* that were not shown in the task. Some participants suggested that the audio security warning should appear earlier in the email list of inbox (see Fig. 7). They mentioned they often read emails directly using the email list because it was easier for screen readers to find important information such as email sender, subject, and even the email body in a clean *table structure*. The email list meets their needs. Therefore, it would be helpful to add a new column in the email list table to include the warning early on. For example, P15 said "*In that (email list) table, you should add a spam column also.*"

Observation 7: Participants have diverse preferences in terms of warning designs. The most preferred design is to use audio warning and present the warning message on the same page of the email body, to reduce the disruption to the email reading.

8. Discussion

8.1. Noticing Warnings and Its Impact

Our result indicates that people with visual impairments tend to miss the existing Gmail warnings. This is because the current warning designs (warning position and HTML tags) do not match with screen reader users' reading habits. Across both the HTML and standard views, screen reader users tend to quickly navigate to the email body using different tactics such as pressing shortcut keys to skip (parts of) the email head. Such behavior is well justified for screen reader users because the email head is crowded with UI components (e.g., the reply button, email label button) that

are non-essential to email reading. As a result, Gmail warnings that are located in the email head can be easily missed by screen reader users. This applies to all the warnings in the HTML view and the icon warning on the standard view. Another reason is the warning's HTML tags. Screen reader users are used to relying on certain shortcut keys to navigate to `link` or heading tags on the web page. A warning message without any `link` or heading can be easily missed (e.g., Gmail's icon warning in the HTML view).

Our inclusive design prototype improves the noticeability of email warnings by implementing the audio warning to draw users' attention. By pressing a shortcut key, the warning message overlays on the whole page, which is hard to miss. This design is consistently implemented for both the HTML view and the standard view.

We did not find statistical evidence that noticing warning significantly improved phishing detection (possibly due to small data size). However, from our qualitative analysis, we observed that participants recognized more phishing emails after switching from the Gmail design to the inclusive design (but not the other way around). Our results also indicate that it is also unrealistic to expect the warning to fully mitigate the phishing risk. We observe that about 25% (5/20) of the users in our study still misidentified a phishing email as legitimate even after they noticed the warning. The result is consistent with prior studies on SSL warnings [63]. Our qualitative data shows that some users do not (fully) trust the warning due to the uncertain tone of the warning. This suggests that further improvement is needed in the wording of the warning to reduce uncertainty. Finally, we show that misplaced warnings on legitimate emails can slightly increase false positives (regardless of the warning design). This result has two implications. First, email receivers should try to avoid misplacing warnings on legitimate emails. Second, this may motivate email senders to fix any misconfigurations on their side to avoid such warnings in their emails.

8.2. Design Implications and Recommendations

Support HTML View. Our study shows more participants prefer the HTML view over the standard View. The HTML view is simpler without dynamic UI components, which is easier for people with visual impairments (and screen reader users) to navigate. We advocate that more email providers should maintain an HTML view like Gmail to support people with visual impairments.

Improve the Consistency of Warnings. Our result shows that the inconsistent warning designs (in Gmail) lead to inconsistent warning noticeability and user confusion. First, the warning on the *same email* looks different between the HTML view and the standard view (in terms of warning position, HTML tags, and even the wording of the warning message). Second, even within the same view, the warnings still have inconsistent positions or HTML tags. We believe that it is reasonable to have different *wordings* for the warning message to indicate different levels of phishing risks. However, the *position* of the warning and the HTML tags used should be as consistent as possible to avoid confusion. In addition, the warning of *the same email* should be consistent across the HTML view and the standard view.

Avoid Putting Warnings in Email Head. Email providers should avoid putting any warning messages in the email head. Screen reader users often miss the warnings in the email head as they tend to quickly skip the UI elements in the head to reach the email body. Putting warnings on the email body or creating a warning overlay are better choices.

Include Link and Heading tags in Warning. To make warnings easily reachable by screen reader users, email providers may consider including `link` tags and heading tags in the warning message. These tags can be easily located by frequently used shortcut keys of screen readers.

Use Audio Warnings. 90% (18/20) of the users in our study prefer the use of audio (bell sound) which helps them notice the warnings. In addition, since screen readers may read emails directly from the email list (without opening an email page), email providers should consider bringing the warning to the email list too to ensure warning noticeability.

Trade-off: Warning Effectiveness vs. User Interruption. Email providers should make careful trade-offs between the warning effectiveness and the amount of user interruption introduced. Our inclusive design prototype uses warning overlay over the whole page. The goal is to introduce disruptions to users' email reading process to increase warning effectiveness. The same idea has been used in existing browser warnings on malicious web pages [53] and URL warnings [63]. However, the interruption can be perceived as too excessive. The alternative design is to include the warning in the email body. While this design still requires screen reader users to navigate to the warning with a sequence of shortcut keys, it does not block the email reading process and has good warning noticeability (see Section 7.2). Our result indicates more users prefer the low-interruption option. A potential middle ground is to use the warning overlay only for high-risk emails (corresponding to the red Gmail banner) and use warning in the email body for other suspicious emails (corresponding to the yellow banner or icon warnings). Finally, email providers may provide the option for users to configure their warning displaying method. However, such a configuration panel should exclude any low-security choices (e.g., warning in the email head) and set a secure "default."

8.3. Generalizability

Our study was conducted on the Gmail desktop interface. Here, we briefly discuss how the lessons we have learned can be potentially generalized to other email providers and interfaces.

Desktop-Based Email. By comparing the email warnings on the desktop (website) versions of Outlook, Yahoo, Zoho, Protonmail, and Gmail, we found that some of the design problems discovered in Gmail commonly exist across email

providers. Among these five popular email platforms, Gmail and Yahoo offer the HTML view (Yahoo Mail Basic). As we tested, similar to Gmail, Yahoo also has inconsistent warning designs between the standard view and the HTML view. As shown in Fig. 10 (in Appendix), the test email we sent to our own account on Yahoo has a warning in the standard view but no warning in the HTML view.

Recall that the icon warning in the Gmail standard view does not match with screen reader users' reading habits. Participants tended to miss icon warnings since they are located in the email head and have no HTML tags. The same problem exists in Outlook. Outlook also has a question mark on the sender photo as the icon warning, which is also located in the email head. While the five email platforms all used banner warnings, some of the banner warnings do not include the HTML tags such as link or heading (e.g., Yahoo Mail). If screen reader users use a shortcut key to go over an email page, such banner warnings could be missed.

Mobile Email Apps. For mobile email apps, we hypothesize that some of the observed problems in web-based Gmail are also applicable, which will need future work for systematic examinations. Our hypothesis is based on the following observations.

First, phone screen readers such as TalkBack (Android) and VoiceOver (iOS) support various touchscreen gestures (e.g., swipe right to select the next item) allowing users to quickly navigate through content on the phone screen. The fast navigation may cause users to miss the warning.

Second, most mobile email apps (e.g., Gmail, Zoho, Protonmail) have similar warning designs as their desktop version. Take the Gmail app for example, it only offers the standard view which has similar icon and banner warnings as the desktop interface. The icon warning is still a question mark on the sender photo (i.e., the same location as the icon warning on the desktop version). However, instead of reading the full warning message in the desktop version, it only reads aloud a short word "unauthenticated" when the screen reader goes over the icon warning on the mobile app. Such a short warning in the email head could be missed. Furthermore, the location of banner warnings in the Gmail app is the same as the desktop version. However, unlike the banner warnings in the desktop version, the banner warnings in the Gmail app do not have a heading tag which could make it more difficult to locate. We defer the experimental validation of these hypotheses to future work.

8.4. Nuisance and Warning Fatigue

Frequent warnings (especially false alerts) may lead to warning fatigue [63]. However, modern email services such as Gmail do not display warnings frequently. The vast majority of spam/phishing emails are directly filtered by email providers (either placed in the spam folder or directly blocked). Only a small number of "uncertain" emails are allowed to enter the inbox with a warning on them [64]. The low frequency of warnings can help reduce warning fatigue. Moreover, to reduce warning nuisance, our designs (e.g., using audio) were derived from the interactive design sessions with our target users, which were generally well-received by later participants (90%). Third, email providers may further reduce misplaced warnings by improving their detection algorithms. Finally, to support users with and without visual impairments, email providers can allow users to choose one of two default "modes": screen reader mode (e.g., audio warnings) and visual mode (e.g., visual warnings).

8.5. Limitations

Our exploratory study has a number of limitations. *First*, we recruited 41 (21+20) participants with visual impairments to qualitatively examine the challenges users face and the design choices that help to improve email warnings' efficacy. Our results are potentially biased towards Gmail and Chrome users, and may not fully reflect the experiences of all users with visual impairments. Our sample size (41) may be smaller than other usability studies, which is primarily due to the difficulty of recruiting people with visual impairments. However, the sample size is comparable with (and is often bigger than) recent studies on this user population (e.g., 8 people in [56], 14 people in [10]). *Second*, due to the small pool of participants, we designed a *within subject* study where one participant reviewed both the inclusive design and the existing Gmail design (for the same 5 emails). To counter the potential ordering effect, about half of the participants reviewed the Gmail design first, and the other half reviewed the inclusive design first. We only used the first-round results for the quantitative analyses. *Third*, we took a role-playing approach using a Gmail account and a virtual machine set up by us. It is possible that participants might behave more (or less) cautiously compared with those using their own accounts/devices. To counter the potential biases, we prompted participants to interact with emails as they would in real life, and further asked questions about their real-life practices. In addition, we pre-installed the screen reader that participants regularly used and instructed them to configure it to make the environment similar to that of their own device.

9. Conclusion

In this paper, we worked with a group of 41 users with visual impairments to examine the challenges they faced when interacting with phishing emails and phishing warnings. Using Gmail as a target platform, our study revealed a number of problems in the current design that made it difficult for screen reader users to notice the warning. Based on the results, we further introduced new designs to improve warning noticeability and help users recognize phishing attempts. We believe more work is needed in this research area to understand and address the challenges faced by people with visual impairments when using existing security and privacy mechanisms.

Acknowledgments. We thank the anonymous reviewers for their constructive comments and suggestions. This work was supported in part by NSF grants 2030521 and 1652497 as well as a JUMP ARCHES award (P336).

References

[1] "World report on vision," World Health Organization (HWO), 2019, https://www.who.int/publications/i/item/9789241516570

[2] M. Vigo and S. Harper, "Coping tactics employed by visually disabled users on the web," *Int. J. Hum.-Comput. Stud.*, vol. 71, no. 11, p. 1013–1025, 2013.

[3] N. G. Sahib, A. Tombros, and T. Stockman, "A comparative analysis of the information-seeking behavior of visually impaired and sighted searchers," *J. Am. Soc. Inf. Sci. Technol.*, vol. 63, no. 2, 2012.

[4] A. Abdolrahmani and R. Kuber, "Should i trust it when i cannot see it? credibility assessment for blind web users," in *Proc. of ACM SIG-ACCESS*, 2016.

[5] T. Ahmed, P. Shaffer, K. Connelly, D. Crandall, and A. Kapadia, "Addressing physical safety, security, and privacy for people with visual impairments," in *Proc. of SOUPS*, 2016.

[6] T. Ahmed, R. Hoyle, P. Shaffer, K. Connelly, D. Crandall, and A. Kapadia, "Understanding the physical safety, security, and privacy concerns of people with visual impairments," *IEEE Internet Computing*, vol. 21, no. 3, pp. 56–63, 2017.

[7] J. Hayes, S. Kaushik, C. E. Price, and Y. Wang, "Cooperative privacy and security: Learning from people with visual impairments and their allies," in *Proc. of SOUPS*, 2019.

[8] T. Ahmed, R. Hoyle, K. Connelly, D. Crandall, and A. Kapadia, "Privacy concerns and behaviors of people with visual impairments," in *Proc. of CHI*, 2015.

[9] T. Akter, B. Dosono, T. Ahmed, A. Kapadia, and B. Semaan, ""i am uncomfortable sharing what i can't see": Privacy concerns of the visually impaired with camera based assistive applications," in *Proc. of USENIX Security*, 2020.

[10] D. Napoli, K. Baig, S. Maqsood, and S. Chiasson, ""i'm literally just hoping this will Work:'' obstacles blocking the online security and privacy of users with visual disabilities," in *Proc. of SOUPS*, 2021.

[11] Z. Zhang, Z. Zhang, H. Yuan, N. M. Barbosa, S. Das, and Y. Wang, "Webally: Making visual task-based captchastransferable for people with visual impairments," in *Proc. of SOUPS*, 2021.

[12] B. Dosono, J. Hayes, and Y. Wang, ""i'm stuck!": A contextual inquiry of people with visual impairments in authentication," in *Proc. of SOUPS*, 2015.

[13] "Most popular email providers in history," Statistics and Data, 2021, https://statisticsanddata.org/data/most-popular-email-providers-in-history/.

[14] H. Hu and G. Wang, "End-to-end measurements of email spoofing attacks," in *Proc. of USENIX Security*, 2018.

[15] "See gmail in standard or basic html version - gmail help." [Online]. Available: https://support.google.com/mail/answer/15049?hl=en

[16] S. Azenkot, K. Rector, R. Ladner, and J. Wobbrock, "Passchords: secure multi-touch authentication for blind people," in *Proc. of ACM SIGACCESS*, 2012.

[17] R. Kuber and S. Sharma, "Toward tactile authentication for blind users," in *Proc. of ACM SIGACCESS*, 2010.

[18] A. Stangl, K. Shiroma, B. Xie, K. R. Fleischmann, and D. Gurari, "Visual content considered private by people who are blind," in *Proc. of ACM SIGACCESS*, 2020.

[19] T. Ahmed, A. Kapadia, V. Potluri, and M. Swaminathan, "Up to a limit? privacy concerns of bystanders and their willingness to share additional information with visually impaired users of assistive technologies," *IMWUT*, vol. 2, no. 3, pp. 1–27, 2018.

[20] T. Akter, T. Ahmed, A. Kapadia, and S. M. Swaminathan, "Privacy considerations of the visually impaired with camera based assistive technologies: Misrepresentation, impropriety, and fairness," in *Proc. of ACM SIGACCESS*, 2020.

[21] D. Napoli, "Accessible and usable security: Exploring visually impaired users' online security and privacy strategies," Ph.D. dissertation, Carleton University, 2018.

[22] A. Oest, P. Zhang, B. Wardman, E. Nunes, J. Burgis, A. Zand, K. Thomas, A. Doupé, and G.-J. Ahn, "Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale," in *Proc. of USENIX Security*, 2020.

[23] P. Peng, C. Xu, L. Quinn, H. Hu, B. Viswanath, and G. Wang, "What happens after you leak your password: Understanding credential sharing on phishing sites," in *Proc. of Asia CCS*, 2019.

[24] A. Oest, Y. Safaei, P. Zhang, B. Wardman, K. Tyers, Y. Shoshitaishvili, A. Doupé, and G.-J. Ahn, "Phishtime: Continuous longitudinal measurement of the effectiveness of anti-phishing blacklists," in *Proc. of USENIX Security*, 2020.

[25] Q. Cui, G.-V. Jourdan, G. V. Bochmann, R. Couturier, and I.-V. Onut, "Tracking phishing attacks over time," in *Proc. of WWW*, 2017.

[26] X. Han, N. Kheir, and D. Balzarotti, "Phisheye: Live monitoring of sandboxed phishing kits," in *Proc. of CCS*, 2016.

[27] J. Vargas, A. C. Bahnsen, S. Villegas, and D. Ingevaldson, "Knowing your enemies: leveraging data analysis to expose phishing patterns against a major us financial institution," in *Proc. of eCrime*, 2016.

[28] Y. Zhang, S. Egelman, L. Cranor, and J. Hong, "Phinding Phish: Evaluating Anti-Phishing Tools," in *Proc. of NDSS*, 2007.

[29] Y. Zhang, J. I. Hong, and L. F. Cranor, "Cantina: a content-based approach to detecting phishing web sites," in *Proc. of WWW*, 2007.

[30] P. Peng, L. Yang, L. Song, and G. Wang, "Opening the blackbox of virustotal: Analyzing online phishing scan engines," in *Proc. of IMC*, 2019.

[31] B. Liu, C. Lu, Z. Li, Y. Liu, H. Duan, S. Hao, and Z. Zhang, "A reexamination of internationalized domain names: The good, the bad and the ugly," in *Proc. of DSN*, 2018.

[32] V. Le Pochat, T. Van Goethem, and W. Joosen, "Funny accents: Exploring genuine interest in internationalized domain names," in *Proc. of PAM*, 2019.

[33] S. Egelman, L. F. Cranor, and J. Hong, "You've been warned: an empirical study of the effectiveness of web browser phishing warnings," in *Proc. of CHI*, 2008.

[34] C. Thompson, M. Shelton, E. Stark, M. Walker, E. Schechter, and A. P. Felt, "The web's identity crisis: Understanding the effectiveness of website identity indicators," in *Proc. of USENIX Security*, 2019.

[35] E. Lin, S. Greenberg, E. Trotter, D. Ma, and J. Aycock, "Does domain highlighting help people identify phishing sites?" in *Proc. of CHI*, 2011.

[36] J. Reynolds, D. Kumar, Z. Ma, R. Subramanian, M. Wu, M. Shelton, J. Mason, E. Stark, and M. Bailey, "Measuring identity confusion with uniform resource locators," in *Proc. of CHI*, 2020.

[37] A. van der Heijden and L. Allodi, "Cognitive triaging of phishing attacks," in *Proc. of USENIX Security*, 2019.

[38] P. Dewan, A. Kashyap, and P. Kumaraguru, "Analyzing social and stylometric features to identify spear phishing emails," in *Proc. of eCrime*, 2014.

[39] S. Duman, K. Kalkan-Cakmakci, M. Egele, W. K. Robertson, and E. Kirda, "Emailprofiler: Spearphishing filtering with header and stylometric features of emails," in *Proc. of COMPSAC*, 2016.

[40] I. Fette, N. Sadeh, and A. Tomasic, "Learning to detect phishing emails," in *Proc. of WWW*, 2007.

[41] J. Hong, "The state of phishing attacks," *Communications of the ACM*, vol. 55, no. 1, 2012.

[42] D. K. McGrath and M. Gupta, "Behind phishing: An examination of phisher modi operandi," in *Proc. of LEET*, 2008.

[43] P. Prakash, M. Kumar, R. R. Kompella, and M. Gupta, "Phishnet: Predictive blacklisting to detect phishing attacks," in *Proc. of INFOCOM*, 2010.

[44] V. Krammer, "Phishing defense against IDN address spoofing attacks," in *Proc. of PST*, 2006.

[45] P. Kumaraguru, Y. Rhee, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge, "Protecting people from phishing: The design and evaluation of an embedded training email system," in *Proc. of CHI*, 2007.

[46] A. Ramachandran, N. Feamster, and S. Vempala, "Filtering spam with behavioral blacklisting," in *Proc. of CCS*, 2007.

[47] G. Ho, A. Sharma, M. Javed, V. Paxson, and D. Wagner, "Detecting credential spearphishing in enterprise settings," in *Proc. of USENIX Security*, 2017.

[48] J. Chen, V. Paxson, and J. Jiang, "Composition kills: A case study of email sender authentication," in *Proc. of USENIX Security*, 2020.

[49] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong, "Teaching johnny not to fall for phish," *ACM Transactions on Internet Technology (TOIT)*, vol. 10, no. 2, pp. 1–31, 2010.

[50] R. Dhamija, J. D. Tygar, and M. Hearst, "Why phishing works," in *Proc. of CHI*, 2006.

[51] N. Stembert, A. Padmos, M. S. Bargh, S. Choenni, and F. Jansen, "A study of preventing email (spear) phishing by enabling human intelligence," in *Proc. of EISIC*, 2015.

[52] J. S. Downs, M. Holbrook, and L. F. Cranor, "Behavioral response to phishing risk," in *Proc. of eCrime*, 2007.

[53] J. Petelka, Y. Zou, and F. Schaub, "Put your warning where your link is: Improving and evaluating email phishing warnings," in *Proc. of CHI*, 2019.

[54] M. Cooper, Y. Levy, L. Wang, and L. Dringus, "Heads-up! an alert and warning system for phishing emails," *Organizational Cybersecurity Journal: Practice, Process and People*, 2021.

[55] G. Sonowal, K. Kuppusamy, and A. Kumar, "Usability evaluation of active anti-phishing browser extensions for persons with visual impairments," in *Proc. of ICACCS*, 2017.

[56] M. Blythe, H. Petrie, and J. A. Clark, "F for fake: Four studies on how we fall for phish," in *Proc. of CHI*, 2011.

[57] J. Zimmerman, J. Forlizzi, and S. Evenson, "Research through design as a method for interaction design research in hci," in *Proc. of CHI*, 2007.

[58] Z. A. Wen, Z. Lin, R. Chen, and E. Andersen, "What.hack: Engaging anti-phishing training through a role-playing phishing simulation game," in *Proc. of CHI*, 2019.

[59] S. Sheng, M. Holbrook, P. Kumaraguru, L. F. Cranor, and J. Downs, "Who falls for phish?: A demographic analysis of phishing susceptibility and effectiveness of interventions," in *Proc. of CHI*, 2010.

[60] C. Marforio, R. Jayaram Masti, C. Soriente, K. Kostiainen, and S. Čapkun, "Evaluation of personalized security indicators as an anti-phishing mechanism for smartphone applications," in *Proc. of CHI*, 2016.

[61] MillerSmiles, "Phishing scams and spoof emails." [Online]. Available: http://www.millersmiles.co.uk

[62] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative research in psychology*, vol. 3, no. 2, pp. 77–101, 2006, publisher: Taylor & Francis.

[63] D. Akhawe and A. P. Felt, "Alice in warningland: A Large-Scale field study of browser security warning effectiveness," in *Proc. of USENIX Security*, 2013.

[64] J. Petelka, Y. Zou, and F. Schaub, "Put your warning where your link is: Improving and evaluating email phishing warnings," in *Proc. of CHI*, 2019.

[65] S. U. IT, "Recent examples of phishing — university it," May 2021. [Online]. Available: https://uit.stanford.edu/phishing

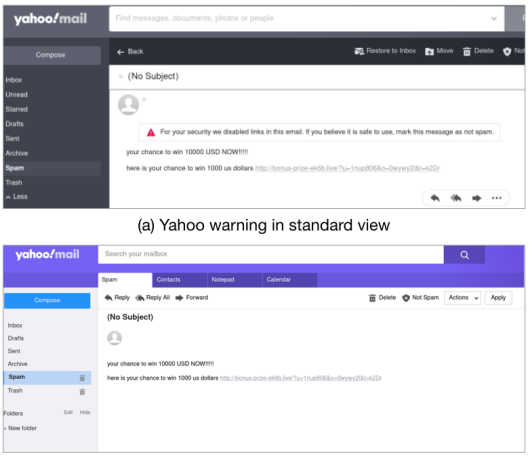

(a) Yahoo warning in standard view

(b) Yahoo warning in basic HTML view

Figure 10: **Inconsistency of email warnings in Yahoo**—The same phishing email has warning in the Yahoo standard view but no warning in the HTML view.

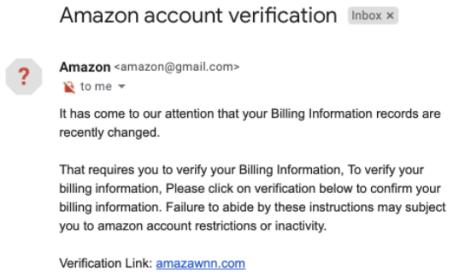

Figure 11: **Example Phishing Email**—Screenshot of one phishing email used in our study. Gmail (standard view) displays a red question mark in the graphical icon as the security warning.

[66] R. E. Boyatzis, *Transforming qualitative information: Thematic analysis and code development*. sage, 1998.

[67] J. L. Fleiss, B. Levin, and M. C. Paik, *Statistical methods for rates and proportions*. john wiley & sons, 2013.

Appendix A. Exploratory Study Task Scenario

Managers or employees in organizations are often overwhelmed with the emails that flood their inboxes. You have been presented with a few emails taken from the inbox of Mr. John King who is one such manager. You work as Mr. John King's assistant. As his assistant, your task is to go through each email present here, to assist Mr. John to process his inbox. Based on your judgment, which of these emails do you think requires Mr. John's attention, and which of these do you think should be ignored? Specifically, what is the action you would take for each of these emails?

Appendix B. Exploratory Study Details

Exploratory Study Protocol. We conducted an interview study with n=21 participants. Due to the pandemic,

Participant ID	Order of Interview	Version of Gmail	Age Group	Gender	Visual Ability	Familiarity with IT/Computer	Primary Occupation
P1	Gmail First	Standard	18-24	Male	Blind	Very Familiar	Computer engineer or IT professional
P2	Inclusive First	HTML	25-34	Male	Blind	Very Familiar	Student
P3	Inclusive First	Standard	25-34	Male	I have some sense of vision at home in physical world. Computer screen is essentially a white screen for me.	Very Familiar	Computer engineer or IT professional
P4	Gmail First	HTML	25-34	Male	I have dark vision and inability to focus sight	Somewhat Familiar	Administrative support (e.g., secretary, assistant)
P5	Inclusive First	Standard	25-34	Male	Complete blindness with light perception	Very Familiar	Business, management, or financial
P6	Inclusive First	HTML	35-44	Female	Totally blind	Somewhat Familiar	Medical (e.g., doctor, nurse, dentist)
P7	Inclusive First	HTML	25-34	Male	Inability to focus the eye lens, blurry vision	Very Familiar	Engineer in other fields
P8	Inclusive First	Standard	18-24	Male	I am a 75% Visually Impaired person. I can't use laptop without any screen reader support.	Very Familiar	Computer engineer or IT professional
P9	Inclusive First	Standard	25-34	Male	Objects looks extremely blurry, I need to take seconds before I can see well	Somewhat Familiar	Business, management, or financial
P10	Gmail First	Standard	55-64	Male	Totally blind	Somewhat Familiar	Self-employed
P11	Gmail First	Standard	18-24	Male	100% legally blind, able to see shadows only if there's lot of sunlight.	Very Familiar	Economics and Commerce
P12	Gmail First	HTML	25-34	Male	I'm visually impaired with more than 90% loss in my vision...	Very Familiar	Computer engineer or IT professional
P13	Gmail First	HTML	35-44	Male	100% blindness in both eyes	Somewhat Familiar	Education (e.g., teacher)
P14	Inclusive First	HTML	35-44	Female	100% blind	Somewhat Familiar	Business, management, or financial
P15	Inclusive First	HTML	25-34	Male	100% blindness	Very Familiar	Education (e.g., teacher)
P16	Inclusive First	HTML	18-24	Male	Blind	Somewhat Familiar	Student
P17	Inclusive First	HTML	35-44	Female	Totally blind	Very Familiar	Computer engineer or IT professional
P18	Gmail First	HTML	45-54	Male	Totally blind	Somewhat Familiar	Education (e.g., teacher)
P19	Gmail First	HTML	25-34	Male	Zero	Very Familiar	Computer engineer or IT professional
P20	Gmail First	Standard	18-24	Male	I'm totally blind	Somewhat Familiar	Student

TABLE 3: Participant demographics (main study).

the interviews were conducted over teleconference software (Zoom). During the study, participants were instructed to read through three emails and determine the action they would take for each one. More specifically, participants logged in to a Gmail account set up by us that contained one legitimate and two phishing emails. The legitimate email was from "Barnes and Noble" (an online bookstore). The first phishing email was selected from the `MillerSmiles.co.uk` scam archive [61]. This email impersonated Amazon as the sender. The second phishing email was selected from Stanford University's public phishing email archive [65]. This email impersonated Microsoft Support Center as the sender. Further details of the three emails are provided in Fig. 13.

Both phishing emails triggered the security warning on the standard view interface of Gmail. Fig. 11 shows one of the phishing emails. The sender icon displays as a red question mark. Hovering over the icon will pop up a short message: "*Gmail couldn't verify [sender address] actually sent this message (and not a spammer).*"

The study procedure worked as the following. Participants first completed a short screening survey (where they self-described their visual abilities)[4]. Then we contacted them to schedule a teleconferencing study session. Study sessions were recorded for later analysis under the consent of participants. During the session, participants played the role of an assistant to help their manager to process and review emails (see the detailed description in Appendix A). By default, participants were logged into the Gmail standard view. During the study, some participants stated that they would prefer switching to the HTML view, and we allowed the switching. Participants opened each email, thinking aloud the action and processes they would take for each email. After reading the three emails, if the participants did not bring up phishing (or talk about the legitimacy of the email) on their own, the interviewer would specifically ask about their opinion of the three emails with regards to phishing. The interviewer would then ask the participant exit interview questions. These questions were related to participants' prior experience with phishing, and their strategies and tools used to detect phishing.

Exploratory Study Results. We recruited participants from social media, mailing lists, and snowball sampling (i.e., participants introduced their friends to join our study). The study took about 1.5 hours (including the setup time). Each participant who completed the study was compensated $30. Also, we compensated additional $10 if they refer a new participant to join the study. Among the n=21 participants, 11 reported female and ten reported male. Nine participants were blind, and the remaining 12 had low vision. All par-

4. Detailed screening questions and exit interview questions for all the studies are shared at https://github.com/yutouzai/inclusive_warning

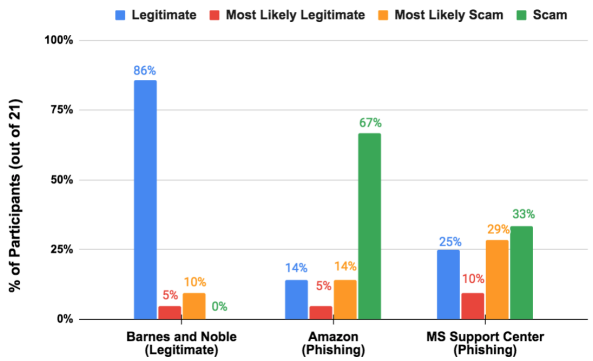

Figure 12: **Phishing Detection Result**—Participants' identification results over the three emails.

ticipants regularly used screen readers.

Inconsistent Warning: HTML vs. Standard Views. During the study, 12 out of the 21 participants switched to the HTML view after they requested it. However, after switching to the HTML view, we observed that there were no longer warnings on the two phishing emails. For the remaining 9 participants who used the standard view, the icon warnings were still shown on the two phishing emails. This indicates that the warnings between the HTML view and the standard view are not consistent.

Phishing Detection. After reading the emails with their screen reader, participants provided assessments on each email regarding their legitimacy. The responses were categorized into four main categories: "Legitimate," "Most Likely Legitimate," "Most Likely Scam," and "Scam." As shown in Fig. 12, most participants correctly recognized the legitimate email from Barnes and Noble (86%). However, many participants still failed to recognize the phishing emails that impersonate Amazon and Microsoft. More participants misclassified the Microsoft phishing email (compared with the Amazon one). Among other reasons, one participant mentioned the Microsoft email did not contain a clickable URL (the email required users to call a specific helpline number) and thus looked legitimate.

Interestingly, the security warning was rarely mentioned by participants as an influencing factor since it was rarely noticed by participants. Among the nine participants that used the Gmail standard view, only two participants noticed at least one warning on their own (one participant noticed the warning only after they were asked about email legitimacy and went back to check the emails again). Based on our observation, most participants either failed to reach the warning or tabbed through the warning quickly without realizing that they had skipped it.

Appendix C. Phishing Detection Results (TN and FN)

The main papers have presented the results of true positives and false positives. Below, we further describe the results of true negatives and false negatives.

True Negatives. True negatives refer to legitimate emails that are correctly identified as legitimate. As shown in Table 1, this counts for the majority of the participants on Email 1 and Email 2. The determination was mostly based on participants' examination of the email content. For example, P13 explained that he still thought Email 2 was legitimate even though it had a warning because the email did not try to collect any information from them. *"They are not asking anything but just saying you will get some books, some audio like that. After replying, they may ask (for) some personal information and I will drop (it) off. But (at) the present it is OK."*

False Negatives. False negatives refer to phishing emails that were incorrectly identified as legitimate by the participants. The most common reason was that the warning on the email was not noticed by the participants. However, sometimes even if the participants noticed the warning, they might still have false negatives (for about 25% of the cases for both Gmail design and inclusive design). One reason is that some participants trusted their own judgments over the warning messages. For example, P10 classified Email 3 as legitimate and explained that warnings could be inaccurate: *"I did see (the warning). I didn't pay attention to that one ... too much obsession with security, so sometimes it is a hindrance to us. So you don't know whether the warning is true or not sometimes."* Another reason is that participants felt the tone of the warning message showed a lack of confidence. For example, on Email 3, the Gmail warning (HTML view) says "*this message may not have been sent by: Google Security do-not-reply@google.secure.com.*" Then participants perceived that Gmail had low confidence in its warning. As mentioned in Section 5, our inclusive design is focused on improving noticeability instead of the wording of the message (we used Gmail's wording). The result indicates that Gmail's current wording can be potentially improved (future work).

Appendix D. Qualitative Analysis Method

All of the interview recordings were transcribed before the data analysis. We conducted a thematic analysis [66], a common method for analyzing qualitative data. One co-author coded each study session, and a second co-author independently coded a subset of sessions. We discussed each other's codes, iterated upon the codes, and finalized a codebook of more than 40 lower-level codes, such as "familiarity with email content," "ignoring promotional emails," and "not care about warnings." The remaining study sessions were coded using this finalized codebook. We then categorized these codes into three high-level themes: "warning noticeability," "phishing detection," and "design preferences." Using the subset of study sessions (40% of the data) coded by the two co-authors independently using the finalized code book, we calculated the inter-coder reliability, which is 0.825 in Cohen's Kappa and is considered good [67].

Email ID	Email Body	Email Sender	Email Subject	Warning [S]	Warning [H]	Phishing
Email 1	Hi, Our community partners with local artists and musician organizations to a summer mini-concert next Sunday from 6 pm to 9 pm. We have every thing from local favorites to Grammy-winning performers right here in the neighborhood park! Feel free to drop by if you are interested. Look forward to seeing you! Thanks, Community Event Team.	event.community@outlook.com	Community Mini Concer	None	None	❌
Email 2	Hi, Welcome to our free audio books list. We are excited to help you discover your next great listen ! Look for our audio book reviews and recommen -dation weekly you'll receive our choices of the best audio books. You'll also occasionally receive special offers! Happy listening. Thanks, Public Library Team	freeaudiobooks.info@gmail.com	Welcome to Free Audiobooks	Gmail couldn't verify that gmail.com actually sent this message (and not a spammer).	Why is this mess age in Spam? It's similar to message that were detected by our spam filters.	❌
Email 3	Dear User, We have detected a new log-in into Google from an iPad device in Naperville, Illinois. If this was not authorized by you, please click here to report https://google.support.com/. Thanks, Google Security	do-not-reply@google.secure.com	Suspicious Account Activity	Be careful with this message. Gmail could not verify that it actually came from do-not-reply@google.secure.com.Avoid clicking links, downloading attachment, or replying with personal information	This message may not have been sent by: Google Security <do-not-reply@google.secure.com> Learn more, Report phishing	✅
Email 4	Hi, You qualify for Apple Music's 3 Month Free Trial. Enjoy 3 months of unlimited music, podcasts, and more. Pay $9.99 /month after the trial ends. Cancel anytime. Click Here to access your free trial today https://apple.music.free.trial.com/. Thanks, Apple Music Team	no-reply@apple.music.com	Apple Music Free Trial	This message seems dangerous. Many people marked similar messages as phishing scams, so this might contain unsafe content. Avoid clicking links, Downloading attachment, or replying with personal information.	Be careful with this message. It conta -ins content that's typically used to steal personal info -rmation. Ignore, I trust this message, Learn more.	✅
Email 5	Hi, You qualify for Hulu's 1 month free trial. Sign -up today, to enjoy one month of unlimited stream -ing of TV shows, movies. Pay 9.99/month after the trial ends. Cancel anytime. Click here to start enjoying what Hulu has to offer https://hulu.free.trial.com. Thanks, Hulu Team	event.community@outlook.com	Community Mini Concert	Gmail couldn't verify that hulu -team.com actually sent this message (and not a spammer)	Why this message in Spam?It's similar to messages that were detected by our spam filters.	✅
Email A	It has come to our attention that your Billing information records are recently changed. That requires you to verify your Billing Information. To verify your billing information, please click on verifi -cation below to confirm your billing information. Failure to abide by these instructions may subject you to amazon account restriction or inactivity. Verification Link: amazawnn.com	amazon@gmail.com	Amazon account verification	Gmail couldn't verify that amazon@gmail.com actually sent this message (and not a spammer)	None	✅
Email B	Microsoft account. Suspicious Account Activity. Hello, This is inform you that we have found suspicious activities with your account. Due to that, we have terminated your windows account. Please find suspicious incident details: Recent Incident Details: Eastern Belarus (IP Address: 10.97.87.25) MAC Address 01:AD:99:00 & IP: 10.97.87.25. If you think this was a mistake and you wish to continue using this windows license key. Please contact technical support at 1-800-341-8835. PS NOTE: Please be at your computer while you call customer support. Windows Help 1-800-341-8835.	2C5A4DD15A5A1106@outlook.com	Account unusual sign-in activity	Gmail couldn't verify that outlook 2C5A4DD15A5A1106@outlook.com actually sent this message (and not a spammer)	None	✅
Email C	Dear John King, Thank you for creating a BN.com account! With your new account you can: View your Order History; Link your Barnes & Noble Membership; Manage your Personal Address Book and Saved Payment Methods; Manage your NOOK Devices and Content; Shop with Instant Purchase; Create Wish Lists; And much, much more. We look forward to your next visit at BN.com. - Barnes & Noble	barnesandnoble@mail.barnesandnoble.com	BARNES & NOBLE	None	None	✅

Figure 13: **Emails and warnings used in the main study and exploratory study**—Email 1-5 are emails in the main study. Email A,B,C are emails in the exploratory study. Warning [S] is the security warning in Gmail standard view. Warning [H] is the security warning in Gmail basic HTML view.